Learning to Transform for Generalizable Instance-wise Invariance

Utkarsh Singhal Carlos Esteves Ameesh Makadia Stella X. Yu

Berkeley UNIVERSITY OF CALIFORNIA · Google Research · UNIVERSITY OF MICHIGAN

Summary: *Given an input image, we predict a distribution over transforms. We use it for data augmentation, aligning instances, and adapting to unexpected (OOD) poses*

Instance-wise transformation distributions

Learning to transform before recognition

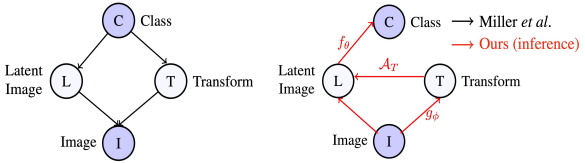

Larger Augmentation Range → Higher Accuracy

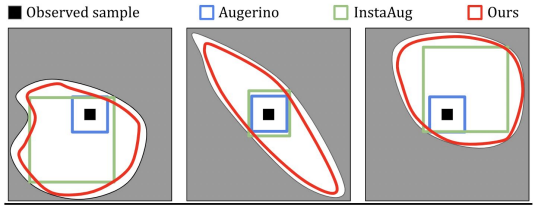

	No Aug.	Fast AA	Augerino	Ours
Acc	90.6	92.65	93.8	**94.3** $\pm$ 0.08

Robustness to novel poses via "mental rotation"

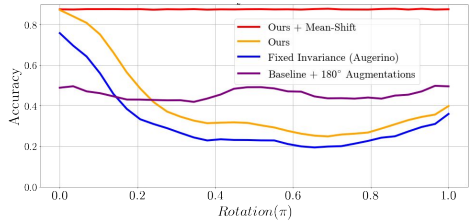

Normalizing flow for instance-wise distribution

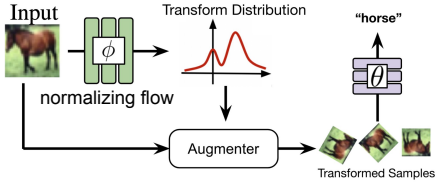

$$\mathcal{L}_{\text{augmenter}}[g_\phi] = \underbrace{\mathcal{L}_{\text{classifier}}[g_\phi]}_{\text{classification}} - \underbrace{\alpha \mathbb{H}[g_\phi]}_{\text{pose range}}$$

Generalizing Invariance across classes

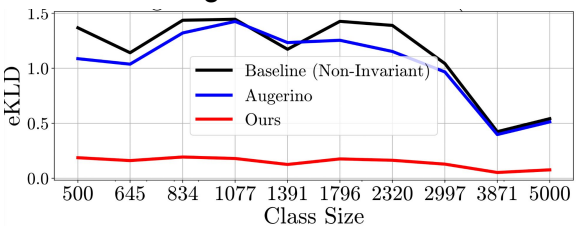

Zero-shot transfer: aligning unseen classes

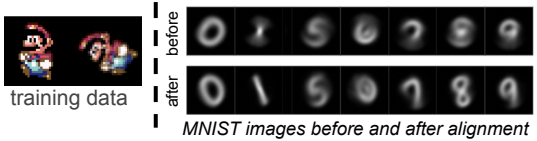

MNIST images before and after alignment

Approximate instance-wise invariance: How?

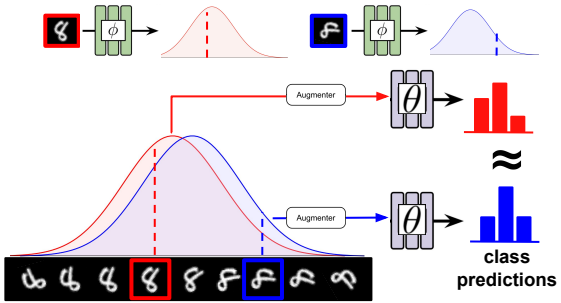

"*Mental rotation*" through mean shift

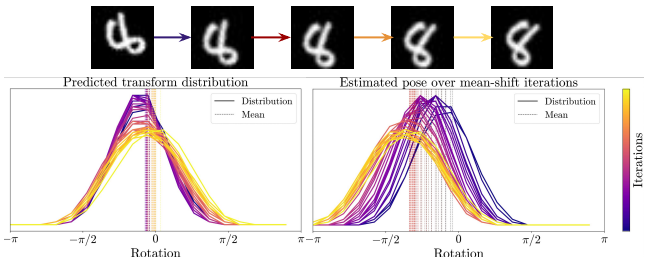

Better imbalanced classification